arXiv:1610.05057v2 [q-bio.PE] 13 Nov 2016

Full Reconstruction of Non-Stationary Strand-Symmetric Models on Rooted Phylogenies

Benjamin D Kaehler

Research School of Biology, Australian National University,

Canberra, Australian Capital Territory, Australia

November 15, 2016

Abstract

Understanding the evolutionary relationship among species is of fundamental importance to the biological sciences. The location of the root in any phylogenetic tree is critical as it gives an order to evolutionary events. None of the popular models of nucleotide evolution used in likelihood or Bayesian methods are able to infer the location of the root without exogenous information. It is known that the most general Markov models of nucleotide substitution can also not identify the location of the root or be fitted to multiple sequence alignments with less than three sequences. We prove that the location of the root and the full model can be identified and statistically consistently estimated for a non-stationary, strand-symmetric substitution model given a multiple sequence alignment with two or more sequences. We also generalise earlier work to provide a practical means of overcoming the computationally intractable problem of labelling hidden states in a phylogenetic model.

Keywords: Identifiability; Phylogenetic inference; Markov process; Maximum likelihood

1 Introduction

The location of the root in a molecular phylogeny has contributed to criminal convictions (González-Candelas et al. 2013), been used to understand the source and epidemiology of human viruses (Podsiadlo and Polz-Dacewicz 2013), determined how biodiversity conservation resources were distributed (Faith and Baker 2006), been used to develop potential HIV vaccines (Nickle et al. 2003), and played an important role in our understanding of the tree of life (Murphy et al. 2007). While an unrooted phylogenetic tree can be used to infer relatedness between species, without the location of the root we know nothing of the order of evolutionary events. It might then be surprising that none of the commonly used Markov models of nucleotide or codon substitution can be used to identify the location of the root without incorporating information that is exogenous to the model. The class of Markov models to which we refer will be made precise in the next section.

ModelTest is one of the most popular pieces of software for selecting phylogenetic models of character substitution (Posada 2008). It allows users to determine which of 88 time-reversible (hereafter *reversible*) substitution models best fits their data. By definition, for a reversible model the location of the root in a phylogeny cannot change the probability distribution of the observed data, the column frequencies. Some software (Knight et al. 2007) allows users to fit non-stationary models of character substitution such as that of Barry and Hartigan (1987a). Unfortunately, the theoretical results that exist around fully general models (Chang 1996) explicitly state that for such models the location of the root is not statistically identifiable, that one cannot use such models to ask where on an edge a root resides, and that it is possible to reformulate the model so that any node is the root.

In practice, the location of the root is usually determined by declaring that a specific taxon in a phylogeny is an *outgroup* or by making a *molecular clock* assumption (Felsenstein 2004). The first method assumes that the location of the root is already known, that it is on the edge connected to the outgroup. The second method comes in varying degrees of complexity. In its most simple form it assumes that the tree is *ultrametric*, that the genetic distance from the root node to each tip is identical. In more sophisticated Bayesian approaches the location of the root enters the calculation as part of the prior distribution of tree topologies and branch lengths, so that the tree is not necessarily ultrametric but that in some sense the evolutionary time from the root of the tree to the tips is the same along every lineage (Drummond and Rambaut 2007).

A third method is to use a substitution model that is able to identify the location of the root. Such a model must not be reversible, but using a model that is not reversible does not automatically ensure that the model is able to identify the root. This statement is easily justified using the findings in Chang (1996), where a model that is non-stationary, so also non-reversible, is not able to recover the root. That a non-reversible model might not be even theoretically able to discover the root seems to have been missed by some authors.

Yang and Roberts (1995) fitted a non-stationary model to rooted topologies of real data using maximum likelihood and found that the location of the root of the tree had a significant effect on likelihood estimates. This is useful empirical evidence but the authors made no attempt to prove that their model is identifiable.

Huelsenbeck et al. (2002) fitted a non-reversible but stationary model to real and simulated data and found that while the outgroup and molecular clock methods were able to recover the location of the root in many cases, their model was not. Again, they made no effort to show that their model is theoretically capable of recovering the location of the root, so the poor performance of their model is not necessarily a reflection on the ability of all substitution models to recover the location of the

root.

Yap and Speed (2005) systematically reproduced the results in Yang and Roberts (1995) and Huelsenbeck et al. (2002) and, with some small discrepancies with the earlier studies, found again that a non-stationary model was able to make statistically significant inferences about the location of the root, but that a non-reversible model did little better than a reversible model. This also was empirical research that left unanswered questions about whether any of the models were theoretically able to identify the location of the root.

The contribution of the present work is to constructively prove that there is a non-stationary substitution process that identifies the location of the root that can be statistically consistently estimated from data. Indeed, the model is shown to be consistently estimable for two taxa. This is not possible for general non-stationary processes (Chang 1996; Bonhomme et al. 2014), so we make the additional assumption that the process is *strand-symmetric*; that the process of evolution is identical on the sense and antisense strands of DNA. The conditions of the proof ensure that the process is non-stationary, and we show that a non-reversible, stationary model is not identifiable so cannot be consistently estimated. This observation sheds some light on the success of non-stationary processes and the failure of non-reversible, stationary processes at detecting the root in the literature.

Much has been written about the biological mechanisms that result in nucleotide substitution processes being strand asymmetric and there is now substantial empirical evidence to support strand asymmetry's existence in nature. Touchon and Rocha (2008) provide a good review of the subject. Strand asymmetry seems to be a localised phenomenon, existing on the scale of genes rather than genomes, and appears to be common in prokaryote and organelle genomes but not in eukaryotes. Nucleotide compositional asymmetry is the most common measure used for statistical inference. Under very loose assumptions (Lobry 1995; Lobry and Lobry 1999) a strand-symmetric process should result in the proportions of As and Ts being equal

and the proportions of Gs and Cs being equal on a single strand. Strand asymmetry can also be inferred by directly comparing estimated rates of nucleotide substitution, although most of the evidence seems to come from counting substitutions in ancestral state reconstructions based on maximum parsimony (eg. Wu and Maeda 1987; Bulmer 1991; Francino and Ochman 2000).

Strand-symmetric models have been used in a maximum likelihood context, although rarely for the purpose of establishing whether strand symmetry is a reasonable assumption. Yap and Pachter (2004) comment that the reversible models that they fitted to real data seemed to exhibit strand symmetry. Squartini and Arndt (2008) fitted a continuous-time, non-stationary, strand-symmetric model on a known, rooted, four-taxon topology to two genome-scale data sets. They comment that their model is not identifiable on the edges incident to the root, but that it is identifiable on the other two edges. This is not strictly correct. As stated in Chang (1996, Remark 2), the labelling of states at the internal nodes is not automatically identifiable even for a continuous-time model. Also, as the mapping from the discrete-time process considered in Chang (1996) to the continuous-time process fitted in Squartini and Arndt (2008) is not always unique (Higham 2008, Section 2.3), continuous-time models are not identifiable under the results in Chang (1996) without further constraints.

Strand-symmetric substitution models have also been approached from a theoretical perspective (Casanellas and Sullivant 2005; Jarvis and Sumner 2013) in the context of reversible models. Jarvis and Sumner (2013) make the observation that strand-symmetric models enjoy the property of *closure*; that a model which is strand-symmetric on two adjacent phylogenetic branches is strand-symmetric across the two branches as well. It is straightforward to show that this property extends to the non-stationary processes that we consider here.

The proofs in this work mirror those in Chang (1996), but apart from adding the assumption of strand symmetry and removing the assumption of an unrooted tree,

we also relax an important assumption that Chang (1996) makes about the structure of the model. As noted in Zou et al. (2011) and addressed in Mossel and Roch (2006) the model of Barry and Hartigan (1987a) is only identifiable up to an arbitrary relabelling of states at internal nodes of the phylogeny. Chang (1996) addresses this problem by assuming that the transition probability matrix for every edge in the topology is *reconstructible from rows*. As we shall demonstrate, this assumption can be restrictive in practice, so, motivated by Remark 4 in Chang (1996), we partially relax it.

In Section 2 we briefly introduce the necessary notation and theoretical context. Section 3 contains the main results of the paper, where we prove that the full topology and parameters of a discrete-time, non-stationary, strand-symmetric Markov model can be recovered from the pairwise joint probability distributions of states between extant taxa. In this section we also extend the result to continuous-time models which are used more commonly in practice. Section 4 proves that the results in Section 3 provide the necessary basis for consistent statistical estimation of the models in question for multiple sequence alignments of increasing length. Section 5 gives some concluding remarks.

2 Markov Models on Trees: Definitions and Notation

We consider a finite set of extant taxa T whose phylogenetic history we wish to infer. The history is modelled as a tree which is a set of nodes S which represent extant or ancestral species and undirected edges E which represent genetic descent. The edges are a set of unordered pairs of nodes, so if $\{r, s\} \in E$, an edge exists between nodes r and s. The nodes consist of the *terminal nodes*, which are just T, and the *internal nodes* N, so that $N = S \setminus T$. The internal nodes represent ancestral taxa

at branching points.

The *degree* of a node is the number of edges incident on that node. We say that a tree is *unrooted* if it contains no internal nodes with degree less than three. A tree is *rooted* if it contains a single node with degree two, which we call the *root*. We do not consider trees with more than one internal node of degree two.

The model of character substitution is properly considered a probabilistic graphical model defined on the tree. That is, we associate a random variable $X_s \in \mathcal{C}$ with each node $s \in S$ so that $\{X_s\}_{s\in S}$ represents the history of a single column in a multiple sequence alignment. Each X_s is independent of all other X_r, conditional on the states at the nodes neighbouring s. As we will focus on the assumption of strand symmetry, we will assume that $\mathcal{C} = \{\mathrm{A}, \mathrm{C}, \mathrm{G}, \mathrm{T}\}$. We also assume that each column in the alignment is an independent observation of the multivariate random variable $\{X_s\}_{s\in T}$.

The model is then specified by a marginal probability row 4-vector π^s, where $\pi^s(i) = \mathrm{P}\left(X_s = i\right)$ for a fixed node s, and a 4×4 transition probability matrix P^{rs} for each edge $\{r, s\} \in E$, where $P^{rs}(i, j) = P(X_s = j | X_r = i)$. As we are interested in rooted tree topologies, it is convenient to characterise the model by π^r where r is the root and the P^{st} where $\{s, t\} \in E$ and t is further than s from r. In this context we will sometimes characterise each P^{st} as the result of a continuous-time Markov substitution process on $\{s, t\}$, in which case $P^{st} = \exp Q^{st}$, where exp is the matrix exponential and Q^{st} is a transition rate matrix.

The statistical challenge is then, for a fixed set of terminal nodes T and a set of n observations of $\{X_s\}_{s\in T}$, which are the n columns in our multiple sequence alignment, to show that as n tends to infinity our estimates of the tree topology and the probabilistic parameters tend to the true values of the generating model.

3 Identifiability of the Full Model

3.1 Identifiability with a Rooted Two-Taxon Topology

It has been known at least since Chang (1996) that a full Markov model is not identifiable for a rooted topology with two terminal taxa without additional assumptions. We shall now prove, under mild conditions, that a non-stationary, strand-symmetric model is identifiable for this topology. This result will provide the foundation for showing that rooted topologies of any size are recoverable, and that the full model is identifiable using joint distributions of states between pairs of taxa.

Definition 3.1. We call a matrix M *strand-symmetric* if it takes the form

$$M = \begin{pmatrix} \delta & \alpha & \beta & \gamma \\ \zeta & \phi & \eta & \theta \\ \theta & \eta & \phi & \zeta \\ \gamma & \beta & \alpha & \delta \end{pmatrix}.$$

Note that if a transition probability matrix is strand-symmetric in this sense, it only fits the usual definition of strand symmetry if the states are properly ordered. Several such orderings are possible. One is (A, C, G, T).

Definition 3.2. We say that a probability 4-vector π is *compositionally asymmetric* if it has non-zero elements and $\pi(1) \neq \pi(4)$, $\pi(2) \neq \pi(3)$, $\pi(1)/\pi(4) \neq \pi(2)/\pi(3)$, and $\pi(1)/\pi(4) \neq \pi(3)/\pi(2)$.

We note that any stationary marginal distribution of a strand-symmetric process violates all four of the conditions of compositional asymmetry, provided that the states are again appropriately ordered as (A, C, G, T) or similar. However, some distributions that are not the stationary distribution of a strand-symmetric process are included in the set of compositionally asymmetric distributions. The reason for this will become apparent in the following proof.

Lemma 3.1. *Take a two-taxon discrete-time Markov model that is defined by a root distribution* $\pi^m = \begin{pmatrix} \pi^m(1) & \dots & \pi^m(4) \end{pmatrix}$ *and two* 4×4 *probability transition matrices* P^{ma} *and* P^{mb}*. Define* $\Pi^m = \operatorname{diag} \pi^m$*. Assume*

(a) P^{ma} *and* P^{mb} *are strand-symmetric;*

(b) π^m *is compositionally asymmetric; and*

(c) P^{ma} *and* P^{mb} *are invertible.*

Then π^m*,* P^{ma}*, and* P^{mb} *are uniquely determined by the joint probability matrix* $J^{ab} = \left(P^{mb}\right)^\intercal \Pi^m P^{ma}$ *up to one of eight permutations of the states at node* m*. That is, the model is identifiable up to a suitable reordering of internal states.*

Example 3.1. It is interesting that a non-reversible, stationary, strand-symmetric model is not identifiable for two taxa, as can be shown by counterexample. Take such a model defined by

$$\pi^m = \begin{pmatrix} 0.20 & 0.30 & 0.30 & 0.20 \end{pmatrix} \quad \text{and}$$

$$P^{ma} = P^{mb} = \begin{pmatrix} 0.38 & 0.25 & 0.21 & 0.16 \\ 0.16 & 0.43 & 0.27 & 0.14 \\ 0.14 & 0.27 & 0.43 & 0.16 \\ 0.16 & 0.21 & 0.25 & 0.38 \end{pmatrix}$$

that yields the joint probability distribution $J^{ab} = \left(P^{mb}\right)^\intercal \Pi P^{ma}$. The model defined

by

$$\pi^m = \begin{pmatrix} 0.20 & 0.30 & 0.30 & 0.20 \end{pmatrix},$$

$$P^{ma} = \begin{pmatrix} 0.35 & 0.34 & 0.12 & 0.19 \\ 0.15 & 0.46 & 0.24 & 0.15 \\ 0.15 & 0.24 & 0.46 & 0.15 \\ 0.19 & 0.12 & 0.34 & 0.35 \end{pmatrix}, \quad \text{and} \quad P^{mb} = \begin{pmatrix} 0.43 & 0.27 & 0.19 & 0.11 \\ 0.07 & 0.39 & 0.31 & 0.23 \\ 0.23 & 0.31 & 0.39 & 0.07 \\ 0.11 & 0.19 & 0.27 & 0.43 \end{pmatrix}$$

is also non-reversible, stationary, and strand-symmetric, and yields the same joint probability distribution J^{ab}. Both models satisfy all constraints of Lemma 3.1 except assumption b. Generating such examples is straightforward, and a Python script for generating random examples is included in the ancillary material. The above example was generated by that script and rounded to two decimal places.

Proof of Lemma 3.1. Single out the permutation matrix

$$S = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix},$$

which has the effect of reversing the order of rows or columns, depending on the direction of multiplication.

First note that if a matrix P is strand-symmetric, then $P = SPS$. Also, $P^\intercal = (SPS)^\intercal = SP^\intercal S$ and $P^{-1} = (SPS)^{-1} = SP^{-1}S$.

By assumptions a and c and because $SS = I$,

$$\begin{aligned} S\left(J^{ab}\right)^{-1} SJ^{ab} &= S\left(P^{ma}\right)^{-1} SS\left(\Pi^m\right)^{-1} \left(\left(P^{mb}\right)^{\intercal}\right)^{-1} S\left(P^{mb}\right)^{\intercal} SS\Pi^m P^{ma} \\ &= \left(P^{ma}\right)^{-1} S\left(\Pi^m\right)^{-1} S\Pi^m P^{ma} \\ &= \left(P^{ma}\right)^{-1} \operatorname{diag}\left(\begin{array}{cccc} \frac{\pi^m(1)}{\pi^m(4)} & \frac{\pi^m(2)}{\pi^m(3)} & \frac{\pi^m(3)}{\pi^m(2)} & \frac{\pi^m(4)}{\pi^m(1)} \end{array}\right) P^{ma}. \end{aligned}$$

The final line is an eigendecomposition of $G = S\left(J^{ab}\right)^{-1} SJ^{ab}$ whose left-eigenvectors are the rows of P^{ma}. The eigenvectors of G are unique by assumption b, so the eigendecomposition is unique up to scaling of the eigenvectors. As the rows of P^{ma} must sum to one, these scaling factors are uniquely determined.

We have shown that the rows of P^{ma} can be uniquely recovered from J^{ab}. The order of these rows is not identifiable without further assumptions, so for the moment we will say that P^{ma} can be identified up to a set of one of eight permutations of the rows such that the resulting P^{ma} is strand-symmetric in form. Then, by assumption c, $\pi^m = \mathbf{1}J^{ab}\left(P^{ma}\right)^{-1}$ and $\left(P^{mb}\right)^{\intercal} = J^{ab}\left(\Pi^m P^{ma}\right)^{-1}$, so π^m and P^{mb} are also identifiable up to one of eight permutations of their elements and rows that correspond to the permutation chosen for P^{ma}. □

3.2 Identifiability of the Topology

In this section we build on a result from Chang (1996) to demonstrate that, under mild conditions, the rooted tree topology is identifiable under a non-stationary, strand-symmetric Markov model from the pairwise distributions of states at terminal nodes.

We first need to also import an equivalence relation between tree topologies. Let $\mathcal{T}_1 = (S_1, E_1)$ and $\mathcal{T}_2 = (S_2, E_2)$ be trees with the same set of terminal nodes T. We say that $\mathcal{T}_1$ and $\mathcal{T}_2$ are *equivalent* if there is a bijective "relabelling" function from S_1 to S_2 that is the identity function for the terminal nodes and such that the edges

E_2 are obtained by applying the function to the nodes in the edges E_1. That is, the topologies $\mathcal{T}_1$ and $\mathcal{T}_2$ are equivalent if they are the same up to a possible relabeling of internal nodes.

The following is Proposition 3.1 in Chang (1996):

Proposition 3.1. *Consider a family of Markov models satisfying the following conditions:*

1. *The edge transition matrices are invertible and not equal to a permutation matrix.*

2. *There is a node v with $\pi^v(i) > 0$ for each $i \in \mathcal{C}$, that is, each character state has positive marginal probability at v.*

Then the unrooted topology is identifiable from the joint distributions of character states at pairs of terminal nodes. That is, if two models in the family induce the same pairwise distributions of character states at their terminal nodes, then the topologies of those two models must be equivalent.

As stated in Chang (1996), Proposition 3.1 follows by combining a result in Buneman (1971) with the *additive* function defined on pairs of nodes as

$$f(\{r, s\}) = -\log\det P^{rs} - \log\det P^{sr}.$$

The function is additive in the sense that $f(\{r, s\}) = \sum_{e \in \text{path}(\{r,s\})} f(e)$, where $\text{path}(\{r, s\})$ is the set of edges joining any two nodes r and s. Chang attributed the $\log\det$ distance to Barry and Hartigan (1987b) and Cavender and Felsenstein (1987). It is unclear whether he was aware that the above definition of f is equivalent to the paralinear genetic distance measure introduced by Lake (1994). In any case we will prefer to formulate this function almost equivalently as

$$f(\{r, s\}) = -\log\det\left(P^{rs}P^{sr}\right), \tag{3.1}$$

because

$$\det\left(P^{rs}P^{sr}\right) = \det\left(\left(\operatorname{diag}\pi^{r}\right)^{-1}\left(P^{sr}\right)^{\intercal}\operatorname{diag}\pi^{s}P^{sr}\right) \geq 0,$$

where π^r and π^s are the marginal probability vectors at r and s respectively, so that the sign of $\det P^{rs}$ becomes irrelevant. The result in Buneman (1971) shows that if the values of an additive function are known between the pairs of all terminal nodes, then under the conditions of Proposition 3.1, the topology and the value of the function on all edges of the topology are determined. As f can be calculated from pairwise distributions between terminal nodes, so the topology and the value of f on every edge can be determined.

Corollary 3.1. *Suppose the evolutionary tree has nodes of degree three or one, with the exception of one special node which we designate the root. The root can have degree one, two, or three. Assume that*

(A) π^m is compositionally asymmetric for every internal node m.

Assume also that for each edge $\{u, v\}$, where v is further from the root than u,

(B) P^{uv} is invertible,

(C) P^{uv} is not a permutation matrix, and

(D) P^{uv} is strand-symmetric.

Then the rooted topology of the tree is recoverable from pairwise distributions of states at terminal nodes. That is, if two models in the family induce the same pairwise distributions of character states at the terminal nodes, then the topologies recovered by those two models must be equivalent.

Proof. Under assumptions A, B, and D, Lemma 3.1 can be applied to any pair of terminal nodes a and b to obtain π^s and P^{sa}, where s is the most recent common

ancestor of a and b, up to a consistent permutation of their elements and rows. We can then calculate $P^{as} = (\operatorname{diag} \pi^a)^{-1} (P^{sa})^\intercal \operatorname{diag} \pi^s$, and so $f(\{s, a\})$. Note that the value of $f(\{s, a\})$ is independent of the chosen permutation of P^{sa} by (3.1).

Now fix a and calculate $c = \arg\max_{b \in T} f^b(\{s, a\})$ where $f^b(\{s, a\})$ is the value of f for a and the most recent common ancestor of a and b, and T are the terminal nodes. The root is then the most recent common ancestor of a and c.

By assumptions A, B, and C, Proposition 3.1 applies and the unrooted topology is identifiable. As for Proposition 3.1, we can also determine the value of f for every edge in the unrooted topology. As the most recent common ancestor of a and c must lie on the path from a to c, $f^c(\{s, a\})$ then gives us the address of the root on that path. That is, we have determined the rooted topology and the value of f for every edge in the topology. □

3.3 Pairwise Distributions Determine the Model

We now approach the main result in Theorem 3.1. We show that a non-stationary, strand-symmetric model is fully identifiable given pairwise joint state distributions between all taxa. Up to this point we have made limited assertions about the labelling of states at internal nodes, but we will now make stronger assumptions to remove this ambiguity. One value of knowing this labelling is that it will allow us to fit continuous-time models.

Start by defining two sets of matrices, the first of which is previously defined in Chang (1996).

Definition 3.3. A set of matrices $\mathcal{M}$ is *reconstructible from rows* if for each $M \in \mathcal{M}$ and each permutation matrix $R \neq I$, we have $RM \notin \mathcal{M}$.

As we mentioned in the introduction, the assumption in Chang (1996) that every transition probability matrix belong to a class of matrices that is reconstructible from rows is used to identify the labelling of states at internal nodes, which otherwise

would not be identifiable.

Example 3.2. To illustrate the abstract concept of a set of matrices that is reconstructible from rows, we borrow an example from (Chang 1996). We say that a square matrix P is *diagonal largest in column* (DLC) if $P(j,j) > P(i,j)$ for all $i \neq j$. The set of DLC matrices is reconstructible from rows.

While DLC serves as a good means for conceptualising a set of matrices that is reconstructible from rows, it has also been shown to be useful as an empirical test for model identifiability (Kaehler et al. 2015). As remarked in Chang (1996), it is not the only set of matrices that is reconstructible from rows but might be a reasonable assumption. Unfortunately, as branch lengths increase, the assumption of DLC becomes less reasonable.

Example 3.3. Under very weak assumptions, as branch length increases the transition probability matrix on that branch must tend towards its stationary limit. That is, all its rows must tend towards equality. Rows that are closer to equality are in some sense less likely to satisfy DLC. To illustrate the point,

$$P = \begin{pmatrix} 0.5 & 0.3 & 0.2 \\ 0.1 & 0.4 & 0.5 \\ 0.1 & 0.1 & 0.8 \end{pmatrix} \quad \text{is DLC.}$$

$$P^2 = \begin{pmatrix} 0.3 & 0.29 & 0.41 \\ 0.14 & 0.24 & 0.62 \\ 0.14 & 0.15 & 0.71 \end{pmatrix} \quad \text{is not DLC.}$$

However, it is an important observation that there is no way of permuting the rows of the second matrix in the above example so that it is DLC. We are therefore motivated to coin the following term.

Definition 3.4. A set of matrices $\mathcal{S}(\mathcal{M})$ is *sympathetic* to a set of matrices $\mathcal{M}$ if

it contains all matrices M such that for every permutation matrix $R \neq I$, we have $RM \notin \mathcal{M}$.

Note that if $\mathcal{M}$ is reconstructible from rows then $\mathcal{M} \subseteq \mathcal{S}(\mathcal{M})$. To slightly abuse the terminology, sympathetic matrices are useful because they provide matrices that might not be reconstructible from rows, but that will not contradict the ordering of states of internal nodes that is suggested by matrices that are.

Example 3.4. The second matrix in Example 3.3 is not DLC but is from the set of matrices that is sympathetic to the set of DLC matrices.

The following theorem extends and restricts Theorem 4.1 in Chang (1996). It makes an additional assumption of strand symmetry, but in so doing is able to accommodate rooted topologies. Motivated by Remark 4 in Chang (1996), it relaxes the assumption that every edge must be reconstructible from rows. Theorem 4.1 in Chang (1996) allows *degenerate* topologies, where nodes are allowed to have degree greater than three. We assume nondegenerate topologies for ease of exposition and because we will only make that assumption later anyway when we prove that the model can be statistically consistently estimated, as Chang does in his Theorem 5.1. We note that extension of the following theorem to degenerate topologies should be straightforward.

Theorem 3.1. *Assume the conditions for Corollary 3.1. Define $\mathcal{M}$ as a set of matrices that is reconstructible from rows and $\mathcal{S}(\mathcal{M})$ as the set of matrices that are sympathetic to $\mathcal{M}$. Further assume that*

(E) there exists a path from every internal node m to a terminal node such that for each edge $\{u, v\}$ in the path, where u is closer to m than v, $P^{uv} \in \mathcal{M}$, and that

(F) $P^{uv} \in \mathcal{S}(\mathcal{M})$ for every edge $\{u, v\}$ in the topology.

Then the full model is identifiable. That is, the topology and all of the transition probability matrices are uniquely determined by the joint distribution of character states at the terminal nodes of the tree.

Remark 3.1. The relaxation of the assumption that every transition probability matrix must be reconstructible from rows to the assumption that every matrix must be sympathetic to those which are might not seem important, but in fact it greatly increases the range of models to which the theorem applies. By allowing sympathetic transition probability matrices, we are in effect allowing long branches in the tree as long as there are sufficient short branches to consistently reconstruct the labelling of states at internal nodes.

Proof of Theorem 3.1. We have from Corollary 3.1 the full rooted topology. It is enough to then determine a transition probability matrix for each edge. We will proceed by induction.

If the topology has two terminal nodes and one node is the root then the full model is trivially determined.

If the topology has two terminal nodes and neither node is the root, then by assumptions A, B, and D we can apply Lemma 3.1 to recover the model up to a permutation of labels of internal states. At least one transition probability matrix must be reconstructible from rows by assumption E and the other at worst will not give an alternative ordering by assumption F, so the ordering of states at the internal node can be recovered and the full model is determined.

If the topology has more than two terminal nodes, choose an arbitrary terminal node a and denote its neighbouring node m. We treat two cases separately.

If a is not the root, choose another terminal node b such that m is the most recent common ancestor of a and b. By assumptions A, B, and D, we can apply Lemma 3.1 to a and b to determine P^{ma} up to a permutation of labels of internal states. By assumption F we can deduce whether $P^{ma} \in \mathcal{M}$, and if it is we can infer

the correct permutation to recover P^{ma}. Otherwise, restart this iteration with a different a.

If a is the root, choose two other terminal nodes b and c such that the most recent common ancestor of b and c is m. Again by assumptions A, B, and D, we can apply Lemma 3.1 to b and c to obtain a stochastic matrix V and a diagonal matrix U such that $P^{mb} = RV$ and $\operatorname{diag} \pi^m = RUR^\intercal$ for some permutation matrix R. Then $P^{ma} = RU^{-1} \left(V^{-1}\right)^\intercal J^{ba}$. Once more if $P^{ma} \in \mathcal{M}$, we infer the correct matrix R and we know P^{ma}. Otherwise, restart this iteration with a different a.

So we are able to determine P^{ma} for a terminal node a and its neighbour m. The induction step is then to remove the edge $\{m, a\}$ from the tree, and make m a terminal node in two subtrees, or one subtree if m happens to be the root. For any subtree, $J^{im} = J^{ia} \left(P^{ma}\right)^{-1}$ for any terminal node of the subtree i. This step can be repeated until all subtrees have two taxa and the full model is recovered on all edges. □

3.4 Continuous-Time Models

We will now show that Theorem 3.1 is still applicable in almost all cases if we further restrict the model to be continuous-time. We state this explicitly as some care is required because the matrix logarithm can have multiple roots. Continuous-time models provide more information than discrete-time models, for instance for questions concerning genetic distance and relative rates of evolution.

Theorem 3.2. *Make the assumptions of Theorem 3.1. Additionally assume that for every $\{u, v\} \in E$, where v is further from the root than u, there exists a unique mapping $Q^{uv} = \log P^{uv}$ where every off-diagonal element of Q^{uv} is non-negative. Also replace assumption D with the assumption that Q^{uv} is strand-symmetric. Then the full model is identifiable from the joint distribution of states at the terminal nodes of the tree.*

Remark 3.2. The restriction that the mapping $\log P^{uv}$ be unique may seem like a barrier to implementation of this model. It is difficult to imagine a sufficiently general parametrisation of Q^{uv} that would guarantee this property. However, empirical evidence suggests that it may be better to beg forgiveness than to ask permission (Kaehler et al. 2015). Software exists for checking whether a 4×4 Markov generator enjoys this property, and it was found that it was rarely necessary to actually enforce this constraint for the data sets considered in Kaehler et al. (2015).

Proof. Again set

$$S = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

Note that if $Q = SQS$ for a square matrix Q, then $P = SPS$ for $P = \exp Q$. This follows from the expansion $P = \sum_{i=1}^{\infty}(n!)^{-1}Q^n$ and that $Q^n = SQ^nS$ because for example $SQ^nS = SQSSQS \ldots SQS$. The assumption that Q^{uv} is strand-symmetric therefore also ensures that assumption D of Theorem 3.1 is satisfied.

By Theorem 3.1, the mapping from the joint distribution of states at the terminal nodes of the tree to the root marginal probability distribution, the tree topology, and the transition probability matrices P^{uv} is unique, so by assumption this property is also enjoyed by the transition rate matrices Q^{uv}. □

4 Reconstruction from Data: Consistency of Maximum Likelihood

As an application of the above identifiability results, we shall now prove that the above non-stationary, strand-symmetric models can be statistically consistently es-

timated from multiple sequence alignments. That is, as the length of the alignment increases, the inferred model must in some sense converge to the correct model.

We will need to use the following definition from Chang (1996), where $\overline{\mathcal{M}}$ denotes the closure of the set of matrices $\mathcal{M}$ under the Euclidean metric.

Definition 4.1. We say that a set of matrices $\mathcal{M}$ is *strongly reconstructible from rows* if, for each $M \in \mathcal{M}$ and each permutation matrix $R \neq I$, we have $RM \notin \overline{\mathcal{M}}$.

We also define the following set, where $B_\epsilon(X)$ is the open ball of radius ϵ centred at X under the Euclidean metric.

Definition 4.2. A set of matrices $\mathcal{S}_\epsilon(\mathcal{M})$ is *strongly sympathetic* to a set of matrices $\mathcal{M}$ if it contains all matrices M such that for every permutation matrix $R \neq I$, we have $B_\epsilon(RM) \cap \mathcal{M} = \varnothing$ for some $\epsilon > 0$.

Theorem 4.1. *Let $\{\mathbb{P}_\theta\}_{\theta\in\Theta}$ be a set of Markov models on trees that have a fixed set of terminal nodes. Suppose the models satisfy the assumptions of Theorem 3.1, with the exception of assumptions E and F, which we restrict slightly to instead assume that from every internal node m there exists a path to a terminal node such that for every edge $\{u, v\}$ in the path, $P^{uv} \in \mathcal{M}$, where v is further from m than u, and for every edge $\{u, v\} \in E$, $P^{uv} \in \mathcal{M} \cup \mathcal{S}_\epsilon(\mathcal{M})$, $\mathcal{M}$ is strongly reconstructible from rows, and $\mathcal{S}_\epsilon(\mathcal{M})$ is strongly sympathetic to $\mathcal{M}$.*

Then the method of maximum likelihood consistently recovers the topology, root marginal probability distribution, and edge transition probability matrices. That is, for any $\theta \in \Theta$, let $\hat{\theta}_n$ denote the maximum likelihood estimate based on n independent observations $\{X_T^i\}_{i\in\{1,\ldots,n\}}$ of character states at the terminal nodes of the tree, then

$$\hat{\theta}_n \xrightarrow{\text{a.s.}} \theta \quad as \quad n \to \infty.$$

Theorem 4.2. *Let $\{\mathbb{P}_\theta\}_{\theta\in\Theta}$ be a set of Markov models on trees that have a fixed set of terminal nodes. Suppose the models satisfy the assumptions of Theorem 4.1.*

Additionally assume that for every $\{u, v\} \in E$, where v is further from the root than u, that there exists a unique mapping $Q^{uv} = \log P^{uv}$, where every off-diagonal element of Q^{uv} is non-negative. Also replace assumption D from Theorem 3.1 with the assumption that $Q^{uv} = SQ^{uv}S$.

Then the method of maximum likelihood consistently recovers the topology, root marginal probability distribution, and edge transition rate matrices. That is, for any $\theta \in \Theta$, let $\hat{\theta}_n$ denote the maximum likelihood estimate based on n independent observations $\{X_T^i\}_{i\in\{1,\ldots,n\}}$ of character states at the terminal nodes of the tree, then $\hat{\theta}_n \xrightarrow{\text{a.s.}} \theta$ as $n \to \infty$.

The following is a restatement of Lemma 5.1 in Chang (1996), where a sketch of its proof is provided.

Lemma 4.1. *Let $\mathcal{X}$ be a finite set and let $\{\mathcal{P}_\theta\}_{\theta\in\Theta}$ be a set of probability distributions on $\mathcal{X}$, where the closure $\overline{\Theta}$ of Θ is a compact subset of a metric space. Let $\{X_i\}_{i\in\mathbb{Z}}$ be independent and identically distributed random variables (or vectors) with probability distribution $\mathcal{P}_{\theta_0}$ for some $\theta_0 \in \Theta$. Assume the identifiability condition*

$$\mathcal{P}_\theta \neq \mathcal{P}_{\theta_0} \quad \text{for each} \quad \theta \in \overline{\Theta} \quad \text{with} \quad \theta \neq \theta_0.$$

Suppose that for each $x \in \mathcal{X}$ the function $\theta \mapsto \mathcal{P}_\theta(x)$ is continuous on $\overline{\Theta}$ and let $\hat{\theta}_n = \hat{\theta}_n(X_1, \ldots, X_n)$ maximise the log likelihood $\sum_{i=1}^n \log \mathcal{P}_\theta(X_i)$ over $\theta \in \overline{\Theta}$. Then under $\mathcal{P}_{\theta_0}$, $\hat{\theta}_n \xrightarrow{\text{a.s.}} \theta_0$ as $n \to \infty$.

Proof of Theorem 4.1. The proof of Theorem 5.1 in Chang (1996) carries to this context with only slight modification. We need to show that the conditions of Lemma 4.1 are implied by the assumptions of Theorem 4.1 with $\mathcal{X} = \{A, C, G, T\}^{|T|}$ and $\mathcal{P}_\theta(A) = \mathbb{P}_\theta(X_T \in A)$ for $A \subseteq \mathcal{X}$.

For $\theta, \theta_0 \in \Theta$, Theorem 3.1 gives that $\mathcal{P}_\theta \neq \mathcal{P}_{\theta_0}$ if $\theta \neq \theta_0$. It remains to show for $\theta_0 \in \Theta$ and $\theta \in \overline{\Theta}$ that $\mathcal{P}_\theta \neq \mathcal{P}_{\theta_0}$ if $\theta \neq \theta_0$.

The proof of Theorem 5.1 in Chang (1996) shows that the assumptions that P^{uv} is invertible and not a permutation matrix for every edge $\{u, v\}$ must be satisfied for a model $\theta \in \overline{\Theta}$ if $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ for some $\theta_0 \in \Theta$, by virtue of the form of $\mathcal{P}_{\theta_0}$. The set of strand-symmetric matrices is closed, so the assumption that each such P^{uv} is strand-symmetric must be satisfied for all $\theta \in \overline{\Theta}$. Finally, if compositional asymmetry is violated at any internal node for a $\theta \in \overline{\Theta}$, then for at least one pair of terminal nodes a and b, the eigendecomposition of $S\left(J^{ab}\right)^{-1} SJ^{ab}$ will have at least one repeated eigenvalue. That is, if $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ for some $\theta_0 \in \Theta$, then for θ we must have that all internal nodes satisfy compositional asymmetry.

So far we have shown that if $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ for $\theta_0 \in \Theta$ and $\theta \in \overline{\Theta}$ then the assumptions of Corollary 3.1 are satisfied by both models and so they must have the same rooted topology.

It remains to show that if $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ for $\theta_0 \in \Theta$ and $\theta \in \overline{\Theta}$ then the transition probability matrices for θ_0 are the same as for θ. We will show that the inductive steps of Theorem 3.1 still apply in this context. As the assumptions of Corollary 3.1 are satisfied for θ, Lemma 3.1 can be safely applied, so the only remaining question is the ordering of rows of the transition probability matrices.

Take a transition probability matrix P from the model $\theta \in \overline{\Theta}$ and the matrix P_0 on the corresponding edge from $\theta_0 \in \Theta$. Firstly assume $P_0 \in \mathcal{M}$. For some permutation matrix R, $P = RP_0$, but if $R \neq I$, $RP_0 \notin \overline{\mathcal{M}}$ as $\mathcal{M}$ is strongly reconstructible from rows. We must also check that $RP_0 \notin \overline{\mathcal{S}_\epsilon(\mathcal{M})}$. By construction, $\overline{\mathcal{S}_\epsilon(\mathcal{M})} \subseteq \mathcal{S}(\mathcal{M})$ and $RP_0 \notin \mathcal{S}(\mathcal{M})$, so $RP_0 \notin \overline{\mathcal{S}_\epsilon(\mathcal{M})}$ and $P = P_0$. Next assume that $P_0 \in \mathcal{S}_\epsilon(\mathcal{M})$. Again $P = RP_0$, but $RP_0 \notin \overline{\mathcal{M}}$. In this case the order of rows of P is decided by the order of those on neighbouring edges, and must be the same for θ and θ_0.

It is therefore possible to modify the proof of Theorem 3.1 such that it ensures that for $\theta_0 \in \Theta$ and $\theta \in \overline{\Theta}$, $\mathcal{P}_\theta \neq \mathcal{P}_{\theta_0}$ if $\theta \neq \theta_0$

We are therefore able to apply Lemma 4.1 and the model can be estimated

consistently under maximum likelihood. □

Proof of Theorem 4.2. This theorem follows from the proofs of Theorems 3.2 and 4.1. The only difficulty is to check whether there might exist a Q which is part of the model θ and a corresponding Q_0 which is a part of the model θ_0 such that $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ but $Q \neq Q_0$. This could only occur if the transition probability matrix P from the same edge and model θ mapped to multiple valid Markov generators. As for the proof of Theorem 4.1, all transition probability matrices are uniquely determined by $\mathcal{P}_\theta$ and as $\mathcal{P}_\theta = \mathcal{P}_{\theta_0}$ we must have that $P = P_0$ for the corresponding P_0 in θ_0, so by assumption Q must be unique and $Q = Q_0$. □

5 Discussion

Assumption A of Corollary 3.1 implies that the process under consideration is non-stationary and we subsequently find that we can recover the root. We further show in example 3.1 that in the stationary, non-reversible case it is not possible to recover the root. This seems to fit with the collective empirical evidence of Yang and Roberts (1995), Huelsenbeck et al. (2002), and Yap and Speed (2005). That is, in some cases a non-stationary process can recover the root, whereas a stationary, non-reversible process cannot. Assumption A is also the source of a subtle and interesting contradiction. It can be shown under mild assumptions that the product of arbitrary stochastic matrices will eventually have almost identical rows as the number of terms in the product increases (Seneta 2006, Section 4.3). It is not difficult to show that the rows must be the stationary distribution of the product, and that if the terms are strand symmetric then the product must be strand symmetric. So for any reasonably long history of strand-symmetric processes, we must have that the marginal probabilities are also strand-symmetric. The conclusion is that for our model to work, the process should probably have been strand asymmetric prior to the root of the tree, and strand symmetric thereafter. An alternative interpretation is that

our model is just a model, and that we could start by assuming that the root probabilities are sufficiently asymmetric for the process to be non-stationary, and then that the process is sufficiently strand-symmetric for the model to be approximately correct. In either case assumption A of Corollary 3.1 could be an interesting tool for probing the limits of inference of this model.

In the spirit of Chang (1996) we have attempted to provide a constructive proof, and hope that the results here presented will enable the implementation of new phylogenetic methods. There are several options for such an implementation. Mossel and Roch (2006) provide an algorithm for directly applying the concepts in Chang (1996) to learn the full model parameters using spectral methods. It would be relatively straightforward to adapt that approach to the proof to Theorem 3.1. Alternatively, progress is being made towards fitting models that are *heterogeneous across lineages* (Jayaswal et al. 2011, 2014). These methods could be adapted to our setting. Also, the method of Yap and Speed (2005) where a homogeneous process is fitted to the whole tree would be trivial to implement for a strand symmetric process, but could be done with a solid theoretical basis and new insight into the limits of inference of such a model.

In addition to these immediate applications, this work provides a foundation for theoretical developments where a non-stationary model can be fully recovered for a two-taxon rooted topology. The ideas presented here could be extended to any such model, not just the strand-symmetric one on which we focus. Another possible direction for future theoretical development is a model that is rate-heterogeneous amongst alignment columns. We have ignored this possibility, both because it is possible to work around it with careful site classification, as in Yap and Speed (2005), and because recent results in non-stationary phylogenetic processes have shown that the general time-reversible model is less biased than the usual rate-heterogeneous general time-reversible model in comparison to a general non-stationary process in some circumstances (Kaehler et al. 2015). Another possible line of enquiry is to re-

fine the sufficient conditions of Theorem 3.1. We know that a stationary process is not identifiable for a rooted topology, but that our non-stationary strand-symmetric process is. Our class of models is slightly narrower than the class of non-stationary, strand-symmetric processes, however, so it may be possible to broaden our sufficient assumptions.

Acknowledgements

I am grateful to Teresa Neeman, Gavin Huttley, Von Bing Yap, Michael Roper, and John Trueman for their helpful comments. I am particularly indebted to Teresa for suggesting Corollary 3.1, which greatly simplified the proof of Theorem 3.1.

This work was supported by the National Health and Medical Research Council [grant number APP1085372].

References

Barry, D. and J. Hartigan (1987a). Statistical analysis of hominoid molecular evolution. Statistical Science 2(2), 191–207.

Barry, D. and J. A. Hartigan (1987b). Asynchronous distance between homologous DNA sequences. Biometrics 43(2), 261–76.

Bonhomme, S., K. Jochmans, and J.-M. Robin (2014). Nonparametric spectral-based estimation of latent structures. Technical report, cemmap working paper, Centre for Microdata Methods and Practice.

Bulmer, M. (1991). Strand symmetry of mutation rates in the β-globin region. Journal of Molecular Evolution 33(4), 305–310.

Buneman, P. (1971). The recovery of trees from measures of dissimilarity. In D. G.

Kendall and P. Tăutu (Eds.), Mathematics in the Archaeological and Historical Sciences, pp. 387–395. Edinburgh University Press, Scotland, UK.

Casanellas, M. and S. Sullivant (2005). The strand symmetric model. In L. Pachter and B. Sturmfels (Eds.), Algebraic Statistics for Computational Biology, Volume 13, pp. 305–321. Cambridge University Press.

Cavender, J. A. and J. Felsenstein (1987). Invariants of phylogenies in a simple case with discrete states. Journal of Classification 4(1), 57–71.

Chang, J. T. (1996). Full reconstruction of Markov models on evolutionary trees: identifiability and consistency. Mathematical Biosciences 137(1), 51–73.

Drummond, A. J. and A. Rambaut (2007). Beast: Bayesian evolutionary analysis by sampling trees. BMC Evolutionary Biology 7(1), 1.

Faith, D. P. and A. M. Baker (2006). Phylogenetic diversity (pd) and biodiversity conservation: some bioinformatics challenges. Evolutionary Bioinformatics 2.

Felsenstein, J. (2004). Inferring Phylogenies (2 ed.). Sinauer Associates Sunderland.

Francino, M. P. and H. Ochman (2000). Strand symmetry around the β-globin origin of replication in primates. Molecular Biology and Evolution 17(3), 416–422.

González-Candelas, F., M. A. Bracho, B. Wróbel, and A. Moya (2013). Molecular evolution in court: analysis of a large hepatitis c virus outbreak from an evolving source. BMC Biology 11(1), 76.

Higham, N. J. (2008). Functions of matrices: theory and computation. Siam.

Huelsenbeck, J. P., J. P. Bollback, and A. M. Levine (2002). Inferring the root of a phylogenetic tree. Systematic Biology 51(1), 32–43.

Jarvis, P. D. and J. G. Sumner (2013). Matrix group structure and Markov invariants in the strand symmetric phylogenetic substitution model. arXiv preprint arXiv:1307.5574.

Jayaswal, V., F. Ababneh, L. S. Jermiin, and J. Robinson (2011). Reducing model complexity of the general Markov model of evolution. Molecular Biology and Evolution 28(11), 3045–3059.

Jayaswal, V., T. K. Wong, J. Robinson, L. Poladian, and L. S. Jermiin (2014). Mixture models of nucleotide sequence evolution that account for heterogeneity in the substitution process across sites and across lineages. Systematic Biology 63(5), 726–742.

Kaehler, B. D., V. B. Yap, R. Zhang, and G. A. Huttley (2015). Genetic distance for a general non-stationary Markov substitution process. Systematic Biology 64(2), 281–293.

Knight, R., P. Maxwell, A. Birmingham, J. Carnes, J. G. Caporaso, B. C. Easton, M. Eaton, M. Hamady, H. Lindsay, Z. Liu, C. Lozupone, D. McDonald, M. Robeson, R. Sammut, S. Smit, M. J. Wakefield, J. Widmann, S. Wikman, S. Wilson, H. Ying, and G. A. Huttley (2007). PyCogent: a toolkit for making sense from sequence. Genome Biology 8(8), R171.

Lake, J. A. (1994). Reconstructing evolutionary trees from DNA and protein sequences: paralinear distances. Proceedings of the National Academy of Sciences 91(4), 1455–1459.

Lobry, J. (1995). Properties of a general model of DNA evolution under no-strand-bias conditions. Journal of Molecular Evolution 40(3), 326–330.

Lobry, J. and C. Lobry (1999). Evolution of DNA base composition under no-strand-

bias conditions when the substitution rates are not constant. Molecular Biology and Evolution 16(6), 719–723.

Mossel, E. and S. Roch (2006). Learning nonsingular phylogenies and hidden Markov models. The Annals of Applied Probability 16(2), 583–614.

Murphy, W. J., T. H. Pringle, T. A. Crider, M. S. Springer, and W. Miller (2007). Using genomic data to unravel the root of the placental mammal phylogeny. Genome Research 17(4), 413–421.

Nickle, D. C., M. A. Jensen, G. S. Gottlieb, D. Shriner, G. H. Learn, A. G. Rodrigo, and J. I. Mullins (2003). Consensus and ancestral state HIV vaccines. Science 299(5612), 1515–1518.

Podsiadlo, L. and M. Polz-Dacewicz (2013). Molecular evolution and phylogenetic implications in clinical research. Annals of Agricultural and Environmental Medicine 20(3).

Posada, D. (2008). jModelTest: phylogenetic model averaging. Molecular Biology and Evolution 25(7), 1253–1256.

Seneta, E. (2006). Non-negative matrices and Markov chains. Springer Science; Business Media.

Squartini, F. and P. F. Arndt (2008). Quantifying the stationarity and time reversibility of the nucleotide substitution process. Molecular Biology and Evolution 25(12), 2525–2535.

Touchon, M. and E. P. Rocha (2008). From GC skews to wavelets: a gentle guide to the analysis of compositional asymmetries in genomic data. Biochimie 90(4), 648–659.

Wu, C.-I. and N. Maeda (1987). Inequality in mutation rates of the two strands of DNA. Nature 327(6118), 169–170.

Yang, Z. and D. Roberts (1995). On the use of nucleic acid sequences to infer early branchings in the tree of life. Molecular Biology and Evolution 12(3), 451–458.

Yap, V. B. and L. Pachter (2004). Identification of evolutionary hotspots in the rodent genomes. Genome Research 14(4), 574–579.

Yap, V. B. and T. Speed (2005). Rooting a phylogenetic tree with nonreversible substitution models. BMC Evolutionary Biology 5, 2.

Zou, L., E. Susko, C. Field, and A. J. Roger (2011). The parameters of the Barry and Hartigan general Markov model are statistically nonidentifiable. Systematic Biology 60(6), 872–5.